

AD-A105 181

RAND CORP. SANTA MONICA CA

F/6 5/10

SCHEMA THEORY AS A GUIDE FOR EDUCATIONAL RESEARCH: WHITE KNIGHT--ETC(U)

MAR 81 P W THORNDYKE

UNCLASSIFIED

RAND/P-6599

NL

1 OF 1
AD 4
10-81B

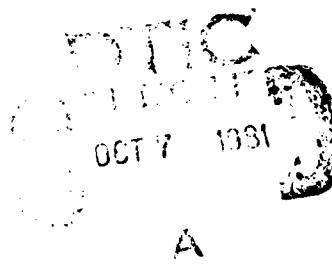
END
DATE
FILED
10-81
OTIC

AD A105181

SCHEMA THEORY AS A GUIDE FOR EDUCATIONAL RESEARCH:
WHITE KNIGHT OR WHITE ELEPHANT?

Perry W. Thorndyke

March 1981



P-6599

8 1 10 7 1981

FILE COPY

The Rand Paper Series

Papers are issued by The Rand Corporation as a service to its professional staff. Their purpose is to facilitate the exchange of ideas among those who share the author's research interests; Papers are not reports prepared in fulfillment of Rand's contracts or grants. Views expressed in a Paper are the author's own, and are not necessarily shared by Rand or its research sponsors.

**The Rand Corporation
Santa Monica, California 90406**

(6) SCHEMA THEORY AS A GUIDE FOR EDUCATIONAL RESEARCH:
WHITE KNIGHT OR WHITE ELEPHANT?

(10) Perry W. Thorndyke

(13) 25

(11) 29 Mar 84

March 29, 1981

(14) RAND/P-6599

Invited address presented at American Educational Research
Association Annual Meeting, Los Angeles, April 1981

296609 x01

SUMMARY

Much recent cognitive and artificial intelligence research has focused on the development of "schema theory." This theory supposes the existence of knowledge structures that encode prototypical descriptions of familiar concepts. Schema theory has developed in a scientific environment that stresses interdisciplinary approaches to the study of intelligent behavior (cognitive science). Consequently, much of the development of schema theory as a descriptive theory of behavior has emphasized its theoretical utility and generality. In contrast, few rigorous tests of the theory have established its psychological validity. Nevertheless, schema theory provides a promising framework for the development of prescriptive methods for learning and reasoning. In particular, schemata may provide (1) memory organizations for use in rapidly acquiring new knowledge, (2) representations for problem-solving strategies, and (3) multi-dimensional data structures for use in problems requiring situation assessment.

A

1. AN OVERVIEW OF SCHEMA THEORY

The cognitive research community has recently witnessed a surge of interest in developing a set of ideas about memory structures that may be collectively referred to as "schema theory." A "schema" comprises a cluster of knowledge that represents a particular generic procedure, object, percept, event, sequence of events, or social situation. This cluster provides a skeleton structure for a concept that can be "instantiated," or filled out, with the detailed properties of the particular instance being represented. For example, a schema for the American Educational Research Association (AERA) conference would encode the standard properties of conferences, such as its location, its date, its type of attendees, its session types, and the length of presentations. In the psychology and computer science literature, knowledge structures encoding prototypical properties of concepts have been variously referred to as frames, scripts, units, objects, as well as schemata. Hereafter, I use the term "schema" to refer to the set of ideas and assumptions common to all of these variants.

By "schema theory," I refer to the collection of models that presume humans encode such knowledge clusters in memory and use them to comprehend and store new instances of the concept. In particular, a schema guides comprehension of new instances of the concept by providing expectations for and constraints on the set of related properties associated with that concept. Thus, if someone asked me to describe the 1982 AERA, I could infer several facts about it, based on my knowledge of the "generic" AERA meeting, without knowing any of the details about the 1982 meetings. For example, I could specify the approximate dates

(the first half of April, beginning on Monday and ending on Friday), type of location (a large U.S. city in a state that has ratified the Equal Rights Amendment), and the nature of the program (a mixture of paper sessions, symposia, and invited addresses) without actually knowing the detailed arrangements.

Numerous papers have recently appeared in the psychological literature interpreting observations or experimental results in terms of the schema theoretic framework. In fact, there has been something of a stampede toward the adoption and development of this notion to explain a broad range of psychological phenomena. In general, most schema-based research has attempted to explain people's ability to comprehend, encode in memory, and recall complex yet familiar aggregations of facts or percepts. Schema theory explains these phenomena by assuming that schemata, or organized collections of facts and relations, are matched against the incoming information and provide a structure in which to encode this information.

This idea is not a new one. Since the turn of the century, various researchers have used essentially the same concept of a schema in such diverse domains as perception, problem solving, cognitive development, and memory for narratives. The significant advance of the recent research over previous work is in the detailed specification of the data structures that encode schematic knowledge. This advance was made possible by the development during the past 20 years of syntactic grammars in linguistics and list processing data structures in computer science.

During the past couple of years, I have found myself becoming increasingly ambivalent in my attitude toward the recent developments in

schema theory. To a great extent, my concerns reflect a set of complex feelings and beliefs about both what the role of an explanatory theory should be in psychology, and what the role of educational and cognitive research should be in our society. On the one hand, experimental psychology has always had a commitment to developing and testing theories of the knowledge structures and mechanisms that support learning, memory, and reasoning. Current work on schema theory, including my own, certainly fits well in that tradition. However, as I shall discuss below, I consider these efforts only partially successful. On the other hand, I believe that it should be the goal of experimental and educational psychologists to translate their theories and laboratory results into prescriptions for human behavior in classroom or everyday situations. This they have been slow to do, particularly regarding the development of pedagogical uses for schemata.

In the remainder of this paper, I shall attempt to defend this prescriptive role of schema theory in educational research. However, to do so, I must place the recent trends in schema theory within a larger context of the cognitive research community and its current directions. Thus, I really have two goals. The first is to argue in behalf of a particular role I believe psychological research should play and a research emphasis that should reflect that role. The second goal is to discuss in detail the particular strengths and weaknesses of schema theory as a research perspective in educational and cognitive psychology. This approach reflects my belief that schema theory provides a metaphor for the style and content of much current psychological research. Therefore, while I focus on schema theory, my comments

reflect in general the pitfalls and opportunities of much current research in psychology.

2. THE CURRENT PSYCHOLOGICAL RESEARCH ENVIRONMENT

One of the major influences on the style and content of recent psychological research has been the evolution of a new discipline called cognitive science. This discipline has created an intellectual zeitgeist within which a new style of research flourishes. Cognitive science research focuses on issues of cognition and intelligence in both humans and computers. These issues include knowledge representation, learning, language and image understanding, inference, spatial cognition, problem solving, and planning. The discipline draws researchers principally from experimental psychology, artificial intelligence (AI), and linguistics.

Two criteria distinguish the research that qualifies as cognitive science from more traditional, discipline-specific work. First, cognitive science seeks to develop theories of complex processing with implications for both the modeling of human performance and the construction of intelligent computer systems. Theories must have broad scope and transcend particular paradigms, even if this is achieved at the expense of analytic granularity. For example, a model of how people scan and compare information in short-term memory is not cognitive science, but how people answer questions could be. Second, there is an explicit emphasis on real, versus artificial, tasks. A model of human learning of nonsense syllables would not be considered cognitive science, while a model of human learning of narrative texts might.

These criteria challenge many of the traditional topics and methods of investigation in experimental psychology and create a tension between the domain of "acceptable" problems and the logical validity of demonstrated results. Many psychologists frequently design experiments with control and treatment groups using simple, artificial laboratory tasks. The conclusions drawn from such experiments may often be valid according to the laws of deductive inference, but the generalizability of the results to normal processing situations is often suspect. On the other hand, tests of more global theories often involve informal methods such as protocol analyses or demonstrations of working computer programs. These methods permit the investigation of more complex processes but do not resolve the uncertain status of the theories as psychologically "valid." This trade-off between problem granularity and theory testability has created an intellectual rift between some psychologists, who view AI system builders as contributing little to psychological modeling, and AI researchers, who view many psychological models as contributing little to the understanding of human behavior.

A second major force shaping psychological research is the socio-political climate in which we work. Pressure appears to be increasing on the cognitive science research community to address more real and applied problems. At The Rand Corporation, an institution whose avowed purpose is to produce research results that impact policies of national importance, I frequently hear the following question from management and my colleagues: "Whose problem will your research solve?" I have also found that during the past five years this question has become increasingly frequent among potential clients in funding agencies. While

university researchers may experience less demand for such "relevance" and "impact," such examinations of research goals are characteristic of a larger national climate that affects us all. As illustrated in Figure 1, several factors seem to be influencing psychologists to consider more applied research questions. First, the wealth of funds for basic research made available by the national Sputnik trauma has slowly diminished as funds have been diverted for military hardware and manpower mobilization. The relative size of this year's proposed cut in the federal budget for social and behavioral sciences research indicates that government views such research as expendable. Second, the availability of basic research jobs in universities has failed to increase as rapidly as the increase in Ph.D.-level researchers, partly due to declines in college enrollment. More and more researchers in psychology have been seeking jobs in applied research institutions and in industry. Third, the basic skills of young adults are slowing declining, and the current technological revolution is creating a demand to train new skills and retrain adults with obsolete skills. This problem is particularly acute in the military, where voluntary enlistment and low wages have reduced personnel quality at a time when the complexity of equipment is increasing rapidly. Thus, there is a need to devise new, high-technology methods for training and instruction. More generally, however, there is a need for psychological research to play a more visible role in maintaining an educated and productive society.

Recently, a colleague of mine, trained as a psychologist but now an AI systems builder, posed the following challenge: "Can general training programs based on cognitive theories contribute as much to these

RESEARCH TRENDS IN EXPERIMENTAL PSYCHOLOGY

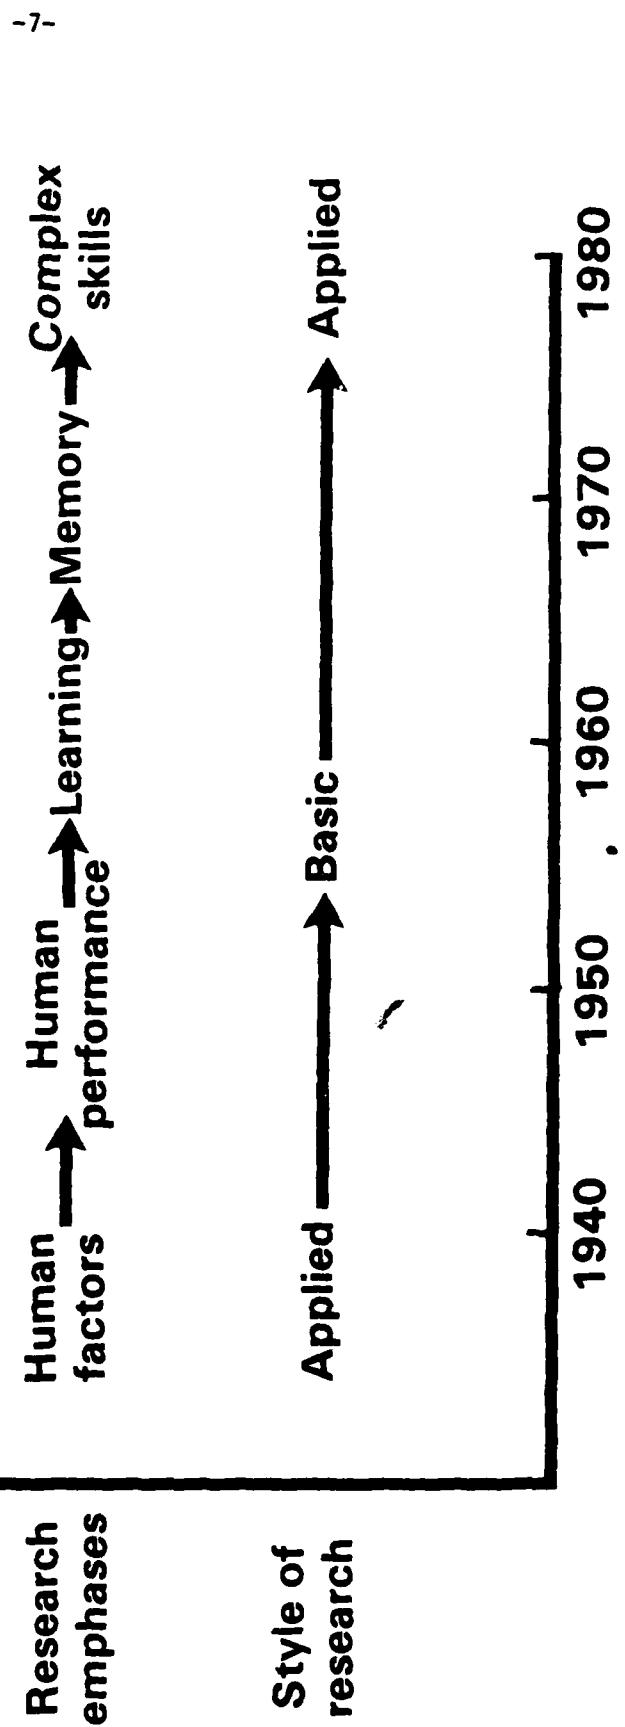


Fig. 1

acknowledged societal needs as domain-specific training?" For example, in Japan, institutions that require Ph.D.-level technical expertise hire students with the equivalent of a B.A. and train job-related skills themselves.

My answer is a qualified yes. Psychologists can contribute to the improvement of human performance. To do so, they must develop robust and general models through constrained theory-building and use these models to devise prescriptive methods for learning and thinking. My affirmative answer is qualified by two observations. First, because of the tension between theory generality and testability that I discussed above, we rarely work seriously at providing the stringent tests required to refine our general theories. These tests are necessary if we are to provide constraint for our theories--that is, qualifications on the conditions under which a theory is valid or useful. Second, we rarely go beyond descriptive theory-building to the specification of prescriptive methods that may follow from the theory. It is precisely such methods that we psychologists should be marketing as the "products" of our trade.

To illustrate these points, I have indicated in Figure 2 a rough sketch of the six-year evolution of schema theory in artificial intelligence and psychology. I take as my starting point the publication in 1975 of Marvin Minsky's paper on frames and the publication of Bobrow and Collins' collection of papers presenting the notion of schemata both as computer data structures and psychological constructs (Minsky, 1975; Bobrow & Collins, 1975). The latter contribution, I might add, contained

THE RECENT EVOLUTION OF SCHEMA THEORY

-9-

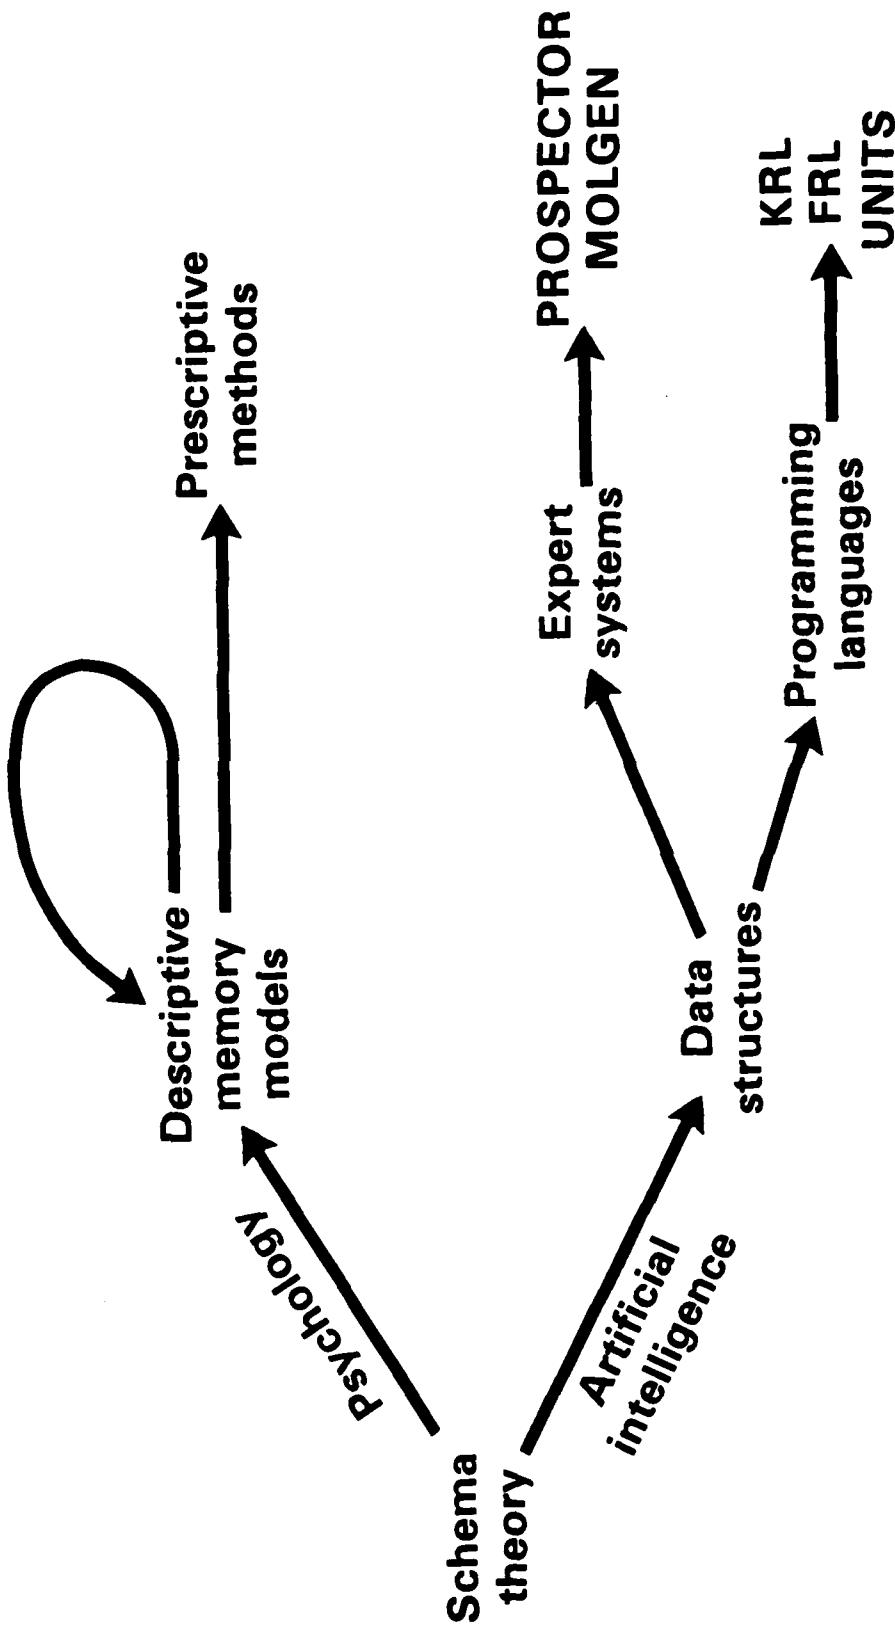


Fig. 2

Rumelhart's seminal work on the use of schemata to represent the structure of narrative texts.

In artificial intelligence research, the concept of memory schemata led to the development of specific data structures to encode schematic knowledge--the AI world's equivalent of memory models. These data structures, in turn, led to the implementation of more general programming languages and supported the development of expert systems to, for example, determine the mineral composition of soil samples (PROSPECTOR) or plan experiments in molecular genetics (MOLGEN). The "proof" of schema theory as a viable formalism for representing knowledge was in the demonstration of working systems--that is, by adding sufficient constraint and detailed specification to the theoretical framework to demonstrate its utility in systems whose performance matched or surpassed that of skilled experts.

Psychologists must also add constraint to their memory models to produce useful systems. Constraint on a psychological model corresponds to theoretical assumptions, empirical predictions, and qualifications on the applicability of the model in explaining learning or thinking in real-world situations. In psychology, this process typically requires experiments that help refine the theory; hence the looping arrow in Figure 2. However, the real payoff of such research comes, as in the AI domain, from the development of "useful systems." For psychologists these systems are the pedagogical applications of the theory for improving people's cognitive behavior. Clearly, the development of these applications must follow progress on the epistemological questions inherent in theory development. However, a shift from the academic

focus to the pragmatic focus may be required for cognitive and educational psychology to survive the coming lean years. And thus far psychologists have done precious little to develop a "cognitive technology" based on principles of schema theory.

I may now state my thesis--my stance regarding schema theory as a guide to educational research--in two parts. First, schema theory is relatively undeveloped as a descriptive psychological theory. In particular, it lacks the constraint necessary to test its viability or to specify the parameters governing its utility in pedagogical situations. Such constrained theory development and testing is critical to the derivation of prescriptive methods based on the theory. Since schema theory has been so highly touted yet so scantily tested, I am inclined to regard it as something of a "white elephant." Second, schema theory provides numerous possibilities for the development of useful techniques for learning, problem solving, and reasoning. If these techniques can be developed and distributed, schema theory may provide cognitive psychology's first "white knight" since the method of loci. In the remainder of this paper, I argue these two points in more detail.

3. DESCRIPTIVE MODELING AND SCHEMA THEORY: THE WHITE ELEPHANT

Since publication of Rumelhart's (1975) application of schema theory to the analysis of the structural dependencies in folk fables, psychologists have proposed a variety of extensions or variations to the basic notion of schemata as prototypical memory structures. Some of these efforts, following Rumelhart's lead, have focused on human learning of narrative texts (e.g., Mandler & Johnson, 1977; Stein & Glenn,

1978; Thorndyke, 1977). Others have developed Schank and Abelson's (1977) notion of scripts as familiar, prototypical event sequences, such as a trip to a restaurant (e.g., Bower, Black, & Turner, 1979). A considerable number of studies have appeared in the literature that purport to support the general, often vague notion of schemata. Typically, this support derives from experimental demonstrations of the validity of schemata as models of memory structures. For example, consider the following sketch of a version of schema theory that might appear in the introduction of a paper:

SCHEMA THEORY, VERSION A

People comprehend and represent narrative texts in memory according to a schema.

The schema guides the hierarchical encoding of knowledge.

Important elements of the overall plot structure appear at the top of the hierarchy.

Less important details, encoding instrumental actions and consequences, appear at the bottom of the hierarchy.

Important information receives more attention at comprehension time and thus has a stronger representation in memory.

PREDICTION B

Therefore, people should recall more high-level information than low-level information.

Experiments typically support the predictions of the theory and thus are taken as evidence for the theory. However, note the fallacious logic of this approach. The argument takes the form:

If A then B.

B.

Therefore, A.

That is, since the prediction is confirmed, the theory is presumably confirmed. In fact, observing B to be true implies nothing about A. If

B is false, however (i.e., the prediction is disconfirmed), then one can conclude that A (the theory) is false. That is, we can only disconfirm, never confirm, theories. When the theory is very carefully constrained and the prediction very specific, we tend to worry less about this problem than when a multitude of other theories might make the same prediction. I fear that the latter situation characterizes many of the predictions based on schema theory.

This situation presents something of a dilemma, however, since as researchers we prefer to develop and defend our favorite theories rather than propose and reject alternatives. In practice, researchers sidestep this dilemma in one of two ways. Both result in the addition of constraint to the theory of interest. The first method involves a process of elimination. Under this method, the researcher designs a set of diagnostic comparisons among alternative theories. The tests are designed so that the competitors make different predictions about a set of observations. A theory whose predictions are confirmed on these observations then receives implicit support relative to the several rejected alternatives, whose predictions are disconfirmed.

The second method of adducing support for a theory involves subjecting the theory to a series of independent tests to which it is vulnerable. These tests, or experiments, take the form:

If A then B1.
If A then B2.
If A then B3.

While none of the confirmed predictions actually confirms A, the repeated accuracy of A's predictions increases its credibility in a

Bayesian sense. That is, the more such tests A passes satisfactorily, the lower the probability that any randomly-selected theory would correctly predict the outcome of all of these tests. In general, the more such predictions A can make, the more constrained A is and the more vulnerable to disconfirmation.

Such systematic refinement through elimination and constraint-building has been somewhat slow to occur in applications of schema theory to psychological modeling. Several factors have contributed to the dearth of these activities. First, most schema theorists have focused primarily on memory structures and relatively little on the processes that utilize these schemata. Thus, there is relatively little constraint on the types of models that could account for observed results. For example, the hierarchical "levels" effect in recall, described above, might occur either because of differential processing of information at storage time, as suggested by Theory A, or because of differential retrieval of information at recall time.

The relative emphasis on structure versus process is, I think, a reflection of the cognitive science zeitgeist in which this theory has developed. Structural models provide the rudiments of a competence model of understanding--one that can be used either for the construction of automated understanding systems or human processing models. Such models thus achieve the interdisciplinary applicability sought by the cognitive science approach. In contrast, the development of process models in psychology entails concern for particularly human activities such as forgetting, retrieval failures and errors, attentional lapses, and so on--in short, a performance model. While mechanisms to account

for these phenomena are crucial to a human processing model, they are less interesting to the designers of intelligent systems, and hence less in the mainstream of cognitive science.

A second deterrent to the critical evaluation of schema theory is that an alternative to schema theory is difficult to formulate. The fundamental concept in schema theory is that people have memory for and predictions about familiar concepts--that is, that they have knowledge. The assumption that knowledge is clustered into related bundles amounts to a prediction about associative strength of related concepts in memory. However, without a definitive statement about what knowledge is or is not part of a schema, or for what concepts schemata do or do not exist, it is impossible to test these strength assumptions.

Finally, schema theory embodies both a popular and a powerful set of ideas, and researchers are understandably more interested in developing their own version of schema theory than in seeking alternatives. We have a tendency to avoid seeking alternative models because (1) it is easier to fine tune an existing theory than to propose a radical alternative, and (2) it is easier to obtain peer support for ideas in a popular area and paradigm than in a radical alternative.

In this regard, the debate over the theoretical adequacy of story schemata between Black and Wilensky (1979) on the one hand, and Rumelhart (1980) and Mandler and Johnson (1980) on the other, published recently in Cognitive Science, is undoubtedly healthy for schema theory. Such debates force schema theorists to defend and refine their models, or even their entire approach, to accommodate criticism and parry attacks. As Kuhn points out in The structure of scientific revolutions

(1962), science progresses through a cyclical process of theory proposal and refinement. A newly proposed theory typically accommodates some set of prior data. Over time, researchers generate new data, both consistent and inconsistent with the theory. The proponents of the theory must patch their theory, by adding constraint, to accommodate the new data. This process continues until the theory is so laden with special cases and qualifications that it is discarded and replaced by an entirely new theory.

For a theory to be subject to this type of evolution, it must possess several properties. I have discussed these in detail elsewhere (Thorndyke & Yekovich, 1980), so I shall only briefly mention them here. In short, a theory must be descriptive, plausible, predictive, and testable. The first two criteria guarantee that the theory can accommodate data that are already available in a more complete, parsimonious, or interesting manner than already existing theories. The third and fourth criteria require that the theory be sufficiently constrained and precise that it predicts the outcome of yet to be conducted experiments and is vulnerable to the results obtained in those experiments. That is, the theory should not be able to accommodate any outcome of these experiments, but should be consistent with only a subset of them. According to these criteria, then, I would not consider psychoanalytic theory a true theory, since no data are in principle inconsistent with the treatment framework. By the same token, the majority of the work in schema theory to date has treated schemata in a rather general, somewhat vague, and imprecise manner, so that it is often difficult to pin down

the precise predictions of the theory and test its assumptions. I do not exclude myself from this criticism.

I confess that I too have been guilty of relatively unconstrained theorizing about schemata. In fact, I proposed something like Theory A above in my early work on human story memory (Thorndyke, 1977, 1978). I have more recently attempted to add some constraint to the schema theoretic framework using both the model elimination methodology and the multiple prediction methodology I outlined above (Thorndyke & Hayes-Roth, 1979; Yekovich & Thorndyke, 1981).

Using the first methodology, Rick Yekovich and I have contrasted several alternative models of how people encode and retrieve information from narrative texts (Yekovich & Thorndyke, 1981). As shown in Table 1, the models were distinguished by their detailed assumptions on each of three issues: (1) whether or not propositions are encoded with differential probabilities as a function of their importance in a narrative structure, (2) whether the representation of the text in memory is hierarchical (as postulated by proponents of story grammars), or heterarchical (as assumed by many associative memory models), and (3) whether memory retrieval depends primarily on direct access to propositions, on a top-down search process, or on a sequential search process. We combined specific assumptions to form eight alternative memory models. For example, Model 2 in Table 1 postulates that subjects use a narrative schema to comprehend stories and encode them in a structured, hierarchical representation. Further, they are more likely to encode information important to the narrative structure than relatively unimportant information. To recall the story later, they retrieve information

from the hierarchical structure using a top-down, depth-first search process. Such a search process will reproduce the correct serial order of the propositions from the text (Thorndyke, 1978).

To evaluate the models, we conducted an experiment in which subjects read and attempted to learn four narrative texts. We then administered both standard recall tests and recognition tests for story propositions. As Table 1 shows, the models make different predictions for the character of the recall and recognition results. Model 2, for example, predicts that high-level information should be both recalled and recognized better than low-level information, reflecting the differential encoding probabilities of these types of information. The presumption of a top-down search through a hierarchical memory representation leads to two other predictions, indicated in columns 7 and 8 of the table. First, the conditional probability of recall of a proposition, given recall of its immediate predecessor ($P(i+1/i)$), should be higher when the two propositions are in the same story constituent (e.g., EPISODE) than when they are in different constituents. This prediction follows from the fact that the number of relations, and hence the search distance, between propositions is in general larger across constituents than within constituents. The second prediction following from the top-down search assumption holds that the conditional probability of recalling a subordinate, or Child, proposition should be higher when its superordinate, or Parent, proposition is recalled than when it is not.

Analysis of the experimental data confirmed both of these predictions. Further, across all stories, propositional recall, but not recognition, varied as a function of importance. These data suggest

Table 1

A SUMMARY OF THE MODELS AND THEIR PREDICTIONS
 (From Yekovich and Thorndyke, 1981)

Model Number	Encoding Bias	Memory Structure	Search Process	Prediction					
				Assumption		Recognition Levels Effect		$P(i+1/i) > P(i/P) > P(C/P) > P(C/\bar{P})$	
				Recall Levels Effect	$P(i+1/i)$	Within	$P(i+1/i)$	Between	$P(i/\bar{P})$
1	Yes	Hierarchical	Direct Access	Yes	Yes	Yes	Yes	No	No
2	Yes	Hierarchical	Top-Down	Yes	Yes	Yes	Yes	Yes	--
3	Yes	Heterarchical	Direct Access	Yes	Yes	Yes	Yes	No	No
4	Yes	Heterarchical	Sequential	Yes	Yes	Yes	Yes	--	Yes
5	No	Hierarchical	Direct Access	No	No	No	No	No	--
6	No	Hierarchical	Top-Down	Yes	No	Yes	Yes	--	--
7	No	Heterarchical	Direct Access	No	No	No	No	No	No
8	No	Heterarchical	Sequential	No	No	No	No	--	Yes

that (1) narrative schemata provide a framework for unbiased encoding of all text propositions, (2) these schemata are hierarchically organized, and (3) the schemata are used for top-down retrieval of information at output time. Thus, with respect to the Models shown in Table 1, only Model 6 accurately predicted all the data.

Studies such as this, that seek to impose constraints on memory models for schemata, are by no means unique in the recent literature (e.g., Kintsch & van Dijk, 1978; Cirilo & Foss, 1980; Chiesi, Spilich & Voss, 1979). However, they represent a minority compared with studies that either extend the purely theoretical assumptions about schemata or test isolated predictions of a particular version of the theory.

As I have argued earlier, the addition of constraint to a theory is pivotal, because these constraints suggest a theory's practical utility. For example, to develop prescriptive methods for the use of schemata in real-world cognition, we require answers to questions such as:

In what types of cognitive activities are schemata useful?
Under what conditions can learned schemata improve performance?
What are the best methods for teaching useful schemata?

In the next section, I propose some tentative answers to these questions and suggest promising directions for the development of prescriptive cognitive methods that utilize schemata.

4. PRESCRIPTIVE METHODS AND SCHEMA THEORY: THE WHITE KNIGHT

I turn now to what I consider to be the more fruitful area for research in schema theory: the development of prescriptive methods for the use of schemata in educational and real-world settings. Schemata

may perform at least three functions in organizing different types of knowledge to support and streamline cognitive processing. They may provide (1) structures for acquiring new knowledge, (2) representations for problem-solving strategies, and (3) multi-dimensional data structures to support situation assessment. I consider each of these functions in more detail below.

Structures for Acquiring Knowledge

As suggested by the research on human story memory, schemata may provide a framework for the organization and encoding of incoming information. When used in this way, schemata provide a sophisticated type of advanced organizer that the understander uses to make sense of and efficiently encode new facts. Unlike some of the psychological research on advanced organizers and contextual cues, however, such a framework provides more than a text title, an explanation of key concepts, or an orienting point of view. Rather, a schema should provide a structural syntax or skeleton on which to "hang" new facts and assertions.

In a study of this application of schemata, Barbara Hayes-Roth and I investigated the utility of previously-learned structural schemata for acquiring new facts (Thorndyke & Hayes-Roth, 1979). Our schemata comprised fairly simple sets of concepts and associations that could be instantiated with different specific facts, as shown in Figure 3. We developed a model within the schema theoretic framework that assumed both costs and benefits for the repeated use of such schemata for the acquisition of new sets of facts. The use of a familiar encoding structure facilitated memory access at storage and retrieval time. However, multiple uses of the shared structure within a short time interval

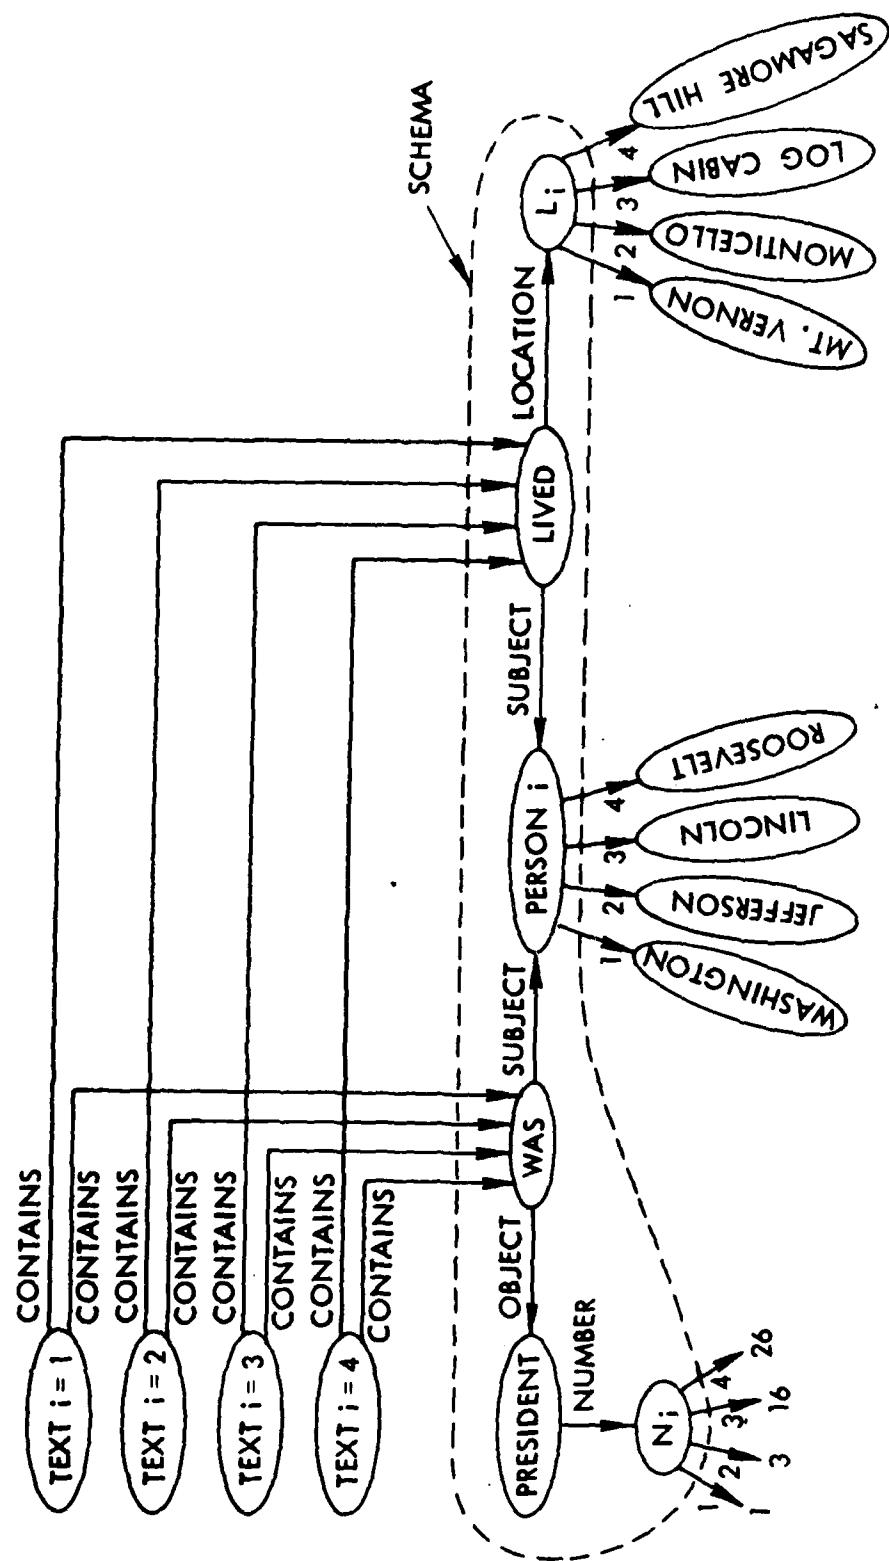


FIG. 3. A shared schema for the representation of the common information.

produced interference among concepts from the various contexts. The combination of these two factors produced a non-monotonic effect on the learnability of new information as a function of the number of prior uses of the schema. To produce an unqualified positive effect of the schema on learnability, it was necessary to space the presentation of fact sets that utilized the schema. This increased the discriminability among the various fact sets and thus eliminated the interference effect. The net result was an unqualified facilitation in learning of new facts that increased with increasing usage of the general schema.

This study illustrates two points germane to the prescriptive use of schemata in natural learning situations. First, while we normally think of schemata as representing very familiar concepts or situations from everyday life, they may also be artificially constructed and taught in order to facilitate the acquisition of new information. Second, the use of familiar memory schemata does not facilitate performance in all situations. Any normative statements about the use of schemata for learning new information must be qualified by the conditions under which they either facilitate or inhibit memory retrieval.

Strategies for Problem-Solving

The concept of the schema has a long history of use in the problem-solving literature. Some of the earliest problem-solving researchers, working between 1910 and 1920, viewed schemata as solution methods or plans of operation guiding the problem solver's behavior (e.g., Woodworth, 1938). Newell and Simon (1972) refined this rather vague notion, though without reference to the term "schema," to suggest a problem-solving method whereby many details of the problem are

initially ignored and the problem is solved at a suitably high level of abstraction. The solution may then be refined by considering the problem details. Such high-level solutions may be thought of as problem-solving schemata because they are general methods that may be instantiated with the details of the particular problem being considered.

I will illustrate this point using the real-world problem-solving domain of air traffic control (Thorndyke, McArthur, & Cammarata, 1981). Frequently, an air traffic controller must make a rapid decision about what flight path an aircraft must take through the air space to avoid collisions with other aircraft in the same area. For example, in Figure 4, "T5" indicates the current location and altitude (5000 feet) of T, and "the "6" in the lower left side of the figure indicates its desired exit point from the airspace. The lower case t's indicate the flight paths resulting from the application of three "strategies" for moving T to 6. Similarly, the symbols "A5" and "U4" indicate the positions and altitudes of two other aircraft in the airspace, and the lower case letters indicate their desired flight paths and destinations (airport "%" for U, airspace exit point "2" for A).

The precise commands required to navigate T along these paths may not be explicitly stored in the controller's memory; indeed, these flight paths may only represent alternative "types" of paths T could take: a westerly routing, a direct routing along the standard airways (the center route), or an easterly routing. Because of the necessity attending immediately to U, an aircraft that wants to be cleared for approach and landing and nearby airport "%," the controller may not have time to generate a complete and detailed plan for T. On the other hand,

a decision must also be made immediately about T since it is destined for a conflict and possible collision with A, which is flying southeast as indicated by the lower case a's. Thus, the controller may select the westerly route as the one that will avoid the conflict with A, and radio T to turn right and head west. Then, making a mental note that he must later transmit a complete plan to T, the controller may then turn his attention to guiding U in to land.

	AIRSPACE DISPLAY	AIRCRAFT DESTINATION	ALT	HDG
N0.....T5.....	T	5	5000 SW
W--E	7 . . t.t, t.t.t.t.t.	A	2	5000 SE
S	. A5 . . , . . , t.t.	U	%	4000 NE
	.t. . . , . . , t. . . t.			
	.t. . a . , , t. . . t.			
	.t. . . *t. . . . t.			
	.t. . . , t, . . . t.			
	.t. . . u-u, t.a, . . t. 4			
	.tU4 . . ut. . . , t. . . . , .			
	3t. . . ut. . . a, t. . . , . .			
	.t. . . ut. . . . t, . , . . .			
	.t. . . %t. . . . t.a			
	.t. . . , t. . . . t,			
	1t. . . , t. . . . , t. . . a . .			
	.t, . . . , t. . . , t.			
	.t. , . . , t. , . . t. a2			
	.t. . . , t, . . . t.			
	.t. . . !t. . . . t.			
	.t. . . , t, . . . t.			
	.t, . t.t, t.t.t, t.			
	6 . t. . . , . . . ,			
, ,			
, ,			
9 , 8			

Fig. 4. Alternative flight paths for aircraft T in an air traffic control sector.

The point of this example is that schematized strategies may streamline problem-solving efforts and reduce processing time. In fact, in simulations of highly time-stressed control situations, we find that experts outperform novices in avoiding aircraft collisions precisely because the experts have a mental library of such high-level strategies from which to develop detailed plans. A major component of acquiring expertise through practice on this control task amounts to learning a set of high-level route descriptions for various combinations of airspace entry points and exit points. The pedagogical value of such schematic problem-solving methods is obvious.

Data Structures for Situation Assessment

People frequently collect data and/or observations from many diverse sources to analyze and build a mental model of the current situation. For example, to decide if it will rain in the morning, one might (1) read the weather forecast in the newspaper, (2) look at the sky, and (3) check an arthritic joint for stiffness. This process may be referred to as situation assessment (Hayes-Roth, 1980).

To decide on the appropriate "model" for the world in such situations, one often relies a conjunction of conditions. If, example, the weather forecast predicts rain, the sky is clear, and my arthritic joint hurts, then I conclude that rain is imminent. Over time, it may be possible to build a relatively accurate method for situation assessment based on a large number of these situation action pairs. The "situation" part of the pair corresponds to a multi-dimensional description of conditions, events, and states of the world that occurred at some point in time. The "action" part of the pair corresponds to the outcome of

interest in the situation (in this example, whether or not it rained and how much).

Such multi-dimensional descriptions may be thought of as schemata since they represent aggregations and abstractions of states and events gleaned from a large number of instances. These descriptions can be matched, or instantiated, with actual observations in a particular new situation in an attempt to select the appropriate model of the situation. (In many cases, the current situation will partially match a number of these schemata; some weighted evaluation function must select the "best" match.)

Such schemata have practical value for both people and systems builders insofar as they provide accurate situation assessments. Two real-world applications of the use of such schemata will serve to illustrate their utility. The first application, "analog" weather modeling, predicts weather by matching a collection of observations of atmospheric and meteorological data against a large data base of previous situation-event pairs (Miller, 1966). For many years, this method was the primary technique for weather prediction, although it is recently been superseded by a more "dynamic," inferential method that relies heavily on the use of thermodynamic and energy exchange equations.

A second application, developed recently at Rand, assesses intentions and likely behaviors of the Soviet Union in a variety of world situations. The data base comprises a catalog of descriptions of various world situations and Soviet responses, including both actual historical data and hypothetical situation-response sets based on Soviet doctrine. Situations are described according to sixteen attributes, thus

defining a sixteen-dimensional space of potential world situations. Each new situation represents a point somewhere in this space. Likely Soviet responses to a new situation may be assessed by determining the distance from its point in the space to the actions associated with nearby points (i.e., other, similar situations). For example, the situation representation surrounding the civil unrest in Afghanistan suggested that a Soviet invasion would be their most likely response. This system provides a realistic tool for foreign policy assessment and strategic analysis, and has promise as a general method for situation assessment.

5. CONCLUSIONS

I have ranged over a large number of issues and examples in this paper in an attempt to defend two somewhat controversial points: (1) the enterprise of cognitive and educational psychology should be the development of cognitive technologies for use in applied settings, and (2) schema theory currently promises more as a prescriptive theory than it delivers as a descriptive theory. To tie these points together, I have argued that we should craft our theories carefully and bring them out of the closet and put them to use as "thinking tools." The decade of the 80's is a particularly critical time to shift our emphasis. We are currently facing a revolution in the amount of knowledge available literally at the fingertips of each person in our society. New advances in computer networking, on-line encyclopedic data bases, sophisticated graphics, and video technologies present tremendous opportunities to train a new generation and a new set of skills. One need only observe

the teenage "Space Invaders" and "Asteroids" aces to appreciate the training and instructional potential of high-technology systems. As educators and cognitive researchers, we should be the leaders in exploiting these tools to guarantee a skilled and knowledge-rich society.

REFERENCES

Bobrow, D., & Collins, A. (Eds.), Representation and understanding. New York: Academic Press, 1975.

Black, J. & Wilensky, R. An evaluation of story grammars. Cognitive Science, 1979, 3, 213-230.

Chiesi, H., Spilich, G., & Voss, J. Acquisition of domain-related information in relation to high and low domain knowledge. Journal of Verbal Learning and Verbal Behavior, 1979, 18, 257-274.

Cirilo, R., & Foss, D. Text structure and reading time for sentences. Journal of Verbal Learning and Verbal Behavior, 1980, 19, 96-109.

Hayes-Roth, F. Artificial intelligence for systems management. In J. Prewitt (Ed.), Machine intelligence 1980: A AAAS symposium. Washington, D.C.: AAAS, 1980.

Mandler, J., & Johnson, N. Remembrance of things parsed: Story structure and recall. Cognitive Psychology, 1977, 9, 111-151.

Mandler, J., & Johnson, N. On throwing out the baby with the bathwater: A reply to Black and Wilensky's evaluation of story grammars. Cognitive Science, 1980, 4, 305-312.

Miller, A. Meteorology. Columbus, Ohio: Merrill Books, 1966.

Minsky, M. A framework for the representation of knowledge. In P. Winston (Ed.), The psychology of computer vision. New York: McGraw-Hill, 1975.

Newell, A., & Simon, H. Human problem solving. Englewood Cliffs, N.J.: Prentice-Hall, 1972.

Rumelhart, D. Notes on a schema for stories. In D. Bobrow & A. Collins (Eds.), Representation and understanding. New York: Academic Press, 1975.

Rumelhart, D. On evaluating story grammars. Cognitive Science, 1980, 4, 313-316.

Schank, R., & Abelson, R. Scripts, plans, and understanding. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1977.

Stein, N., & Glenn, C. An analysis of story comprehension in elementary school children. In R. Freedle (Ed.), Multidisciplinary perspectives in discourse comprehension. Norwood, N.J.: Ablex, 1978.

Thorndyke, P. Cognitive structures in comprehension and memory of narrative discourse. Cognitive Psychology, 1977, 9, 77-110.

Thorndyke, P. Pattern-directed processing of knowledge from texts. In D. Waterman & F. Hayes-Roth (Eds.), Pattern-directed inference systems. New York: Academic Press, 1978.

Thorndyke, P., & Hayes-Roth, B. The use of schemata in the acquisition and transfer of knowledge. Cognitive Psychology, 1979, 11, 82-106.

Thorndyke, P., McArthur, D., & Cammarata, S. AUTOPILOT: A distributed planner for airfleet control. Paper submitted to the International Joint Conference in Artificial Intelligence, 1981.

Thorndyke, P., & Yekovich, F. A critique of schema-based theories of human story memory. Poetics, 1980, 9, 23-49.

Woodworth, R. Experimental psychology. New York: Henry Holt & Co., 1938.

Yekovich, F., & Thorndyke, P. An evaluation of alternative functional models of narrative schemata. Journal of Verbal Learning and Verbal Behavior, 1981, 20, in press.

ATE
LME